

Correlation Hashing Network for Efficient Cross-Modal Retrieval

Yue Cao, Mingsheng Long, Jianmin Wang

School of Software, Tsinghua University, Beijing, China

Tsinghua National Laboratory for Information Science and Technology

caoyue10@gmail.com {mingsheng, jimwang}@tsinghua.edu.cn

Abstract

Due to the storage and retrieval efficiency, hashing has been widely deployed to approximate nearest neighbor search for large-scale multimedia retrieval. Cross-modal hashing, which improves the quality of hash coding by exploiting the semantic correlation across different modalities, has received increasing attention recently. For most existing cross-modal hashing methods, an object is first represented as of vector of hand-crafted or machine-learned features, followed by another separate quantization step that generates binary codes. However, suboptimal hash coding may be produced, because the quantization error is not statistically minimized and the feature representation is not optimally compatible with the binary coding. In this paper, we propose a novel Correlation Hashing Network (CHN) architecture for cross-modal hashing, in which we jointly learn good data representation tailored to hash coding and formally control the quantization error. The CHN model is a hybrid deep architecture constituting four key components: (1) an image network with multiple convolution-pooling layers to extract good image representations, and a text network with several fully-connected layers to extract good text representations; (2) a fully-connected hashing layer to generate modality-specific compact hash codes; (3) a squared cosine loss layer for capturing both cross-modal correlation and within-modal correlation; and (4) a new cosine quantization loss for controlling the quality of the binarized hash codes. Extensive experiments on standard cross-modal retrieval datasets show the proposed CHN model yields substantial boosts over latest state-of-the-art hashing methods.

1. Introduction

While multimedia big data with large volumes and high dimensions are pervasive in search engines and social networks, it has attracted increasing attention to enable approximate nearest neighbors (ANN) search across different media modalities with both computation efficiency and search

quality. As relevant data from different modalities (image and text) may endow semantic correlations, it is important to support *cross-modal retrieval* that returns semantically-relevant results of one modality in response to a query of different modality. An advantageous solution to the cross-modal retrieval is hashing methods [27], which transform high-dimensional data into compact binary codes and generate similar binary codes for similar data items. In this paper, we will focus on cross-modal hashing that builds data-dependent hash coding for efficient cross-media retrieval [21]. Due to large volumes and the semantic gap [23], effective and efficient cross-modal retrieval remains a challenge.

Many cross-modal hashing methods have been proposed to construct the correlation structures across different modalities in the process of hash function learning and indexes cross-modal data into an isomorphic Hamming space [3, 12, 34, 35, 24, 28, 32, 17, 33, 6, 30]. They can be categorized into unsupervised methods [12, 34, 24] and supervised methods [3, 17, 33, 30]. While unsupervised methods are more general and can be trained without semantic labels or relevances, they are restricted by the semantic gap dilemma [23] that high-level semantic description of an object often differs from low-level feature descriptors. Supervised methods can incorporate semantic labels or relevance feedbacks to mitigate the semantic gap [23] and improve the hashing quality, i.e. achieve accurate search with shorter codes.

Recently, deep hashing methods [31, 13] have shown that both feature representation and hash coding can be learned more effectively using deep neural networks [11, 15], which can naturally encode nonlinear hashing functions. Other cross-modal retrieval models via deep learning [19, 25, 28, 6, 26] have shown that deep models can capture nonlinear cross-modal correlations more effectively. These deep hashing methods have created state-of-the-art results on many benchmarks. However, a crucial disadvantage of these cross-modal deep hashing methods is that the quantization error is not statistically minimized hence the feature representation is not optimally compatible with the binary hash coding. Another potential limitation is that they generally do not adopt principled pairwise loss function to link

the pairwise Hamming distances with the pairwise similarity labels and to close the gap between the Hamming distance on binary codes and the metric distance on continuous representations. Therefore, suboptimal hash coding may be produced by existing cross-modal deep hashing methods.

In this paper, we propose Correlation Hashing Network (CHN), a new hybrid architecture for cross-modal hashing. We jointly learn good image and text representations tailored to hash coding and formally control the quantization error. The CHN model constitutes four key components: (1) an image network with multiple convolution-pooling layers to extract good image representations, and a text network with several fully-connected layers to extract good text representations; (2) a fully-connected hashing layer to generate modality-specific compact hash codes; (3) a squared cosine loss layer for capturing both cross-modal correlation and within-modal correlation; and (4) a new cosine quantization loss for controlling the quality of the binarized hash codes. Extensive experiments on standard cross-modal retrieval datasets show the proposed CHN model yields substantial boosts over latest state-of-the-art hashing methods.

2. Related Work

Cross-modal hashing has been an increasingly popular research topic in machine learning, computer vision, and multimedia retrieval [3, 12, 34, 35, 24, 28, 32, 6, 8, 33, 17]. We refer readers to [27] for a comprehensive survey.

Existing cross-modal hashing methods can be roughly categorized into unsupervised methods and supervised methods. IMH [24] and CVH [12] are unsupervised methods, which extend spectral hashing [29] to the multimodal scenarios. CMSSH [3], SCM [33] and QCH [30] are supervised methods, which require that if two points are known to be similar, then their corresponding hash codes from different modalities should be similar. Since supervised methods can explore the semantic labels to enhance the cross-modal correlations and reduce the semantic gap [23], they can achieve superior accuracy than unsupervised methods for cross-modal similarity search with shorter hash codes.

Most of previous cross-modal hashing methods based on shallow architectures cannot exploit *nonlinear* correlation across different modalities. And the latest cross-modal deep models [19, 25, 28, 6] have shown that deep models can capture nonlinear cross-modal correlations more effectively. However, it remains unclear how to maximize cross-modal correlation and control quantization error in a hybrid deep architecture using well-specified loss functions.

In this work, we further extend existing deep hashing methods [31, 13] to cross-modal retrieval by exploiting two key problems: (1) control the quantization error in a principled way, and (2) devise a more principled loss to link the pairwise Hamming distances with the pairwise similarity labels and to reduce the gap between Hamming distance and

cosine distance. These crucial improvements constitute the proposed Correlation Hashing Network (CHN) approach.

3. Correlation Hashing Network

In cross-modal retrieval, the database consists of objects from one modality and the query consists of objects from another modality. We uncover the correlation structure underlying different modalities by learning from a training set of n bimodal objects $\{\mathbf{o}_i = (\mathbf{x}_i, \mathbf{y}_i)\}_{i=1}^n$, where $\mathbf{x}_i \in \mathbb{R}^{d_x}$ denotes the d_x -dimensional feature vector of the image modality, and $\mathbf{y}_i \in \mathbb{R}^{d_y}$ denotes the d_y -dimensional feature vectors of the text modality, respectively. Some pairs of the bimodal objects are associated with similarity labels s_{ij} , where $s_{ij} = 1$ implies \mathbf{o}_i and \mathbf{o}_j are similar and $s_{ij} = -1$ indicates \mathbf{o}_i and \mathbf{o}_j are dissimilar. In supervised hashing, $\mathcal{S} = \{s_{ij}\}$ is usually constructed from the semantic labels of data points or the relevance feedback from click-through data. The goal of CHN is to jointly learn two modality-specific hashing functions $f_x(\mathbf{x}) : \mathbb{R}^{d_x} \mapsto \{-1, 1\}^b$ and $f_y(\mathbf{y}) : \mathbb{R}^{d_y} \mapsto \{-1, 1\}^b$ which respectively encode each unimodal point \mathbf{x} and \mathbf{y} in compact b -bit hash code $\mathbf{h}_x = f_x(\mathbf{x})$ and $\mathbf{h}_y = f_y(\mathbf{y})$ such that the similarity information conveyed in the given bimodal object pairs \mathcal{S} is preserved.

The proposed Correlation Hashing Network (CHN) is a deep architecture for supervised learning to hash, shown in Figure 1. The architecture accepts input in a pairwise form $(\mathbf{o}_i, \mathbf{o}_j, s_{ij})$ and processes them through the deep representation learning and hash coding pipeline: (1) an image network with multiple convolution-pooling layers to extract good image representations, and a text network with several fully-connected layers to extract good text representations; (2) a fully-connected hashing layer to generate modality-specific compact hash codes; (3) a squared cosine loss layer for capturing both cross-modal correlation and within-modal correlation; and (4) a new cosine quantization loss for controlling the quality of the binarized hash codes.

3.1. Hybrid Deep Architecture

Cross-modal data have significantly different statistical properties, which makes it very challenging to capture the correlation structures across modalities based on low-level features. Recently, it has been witnessed that deep learning methods [2], such as deep convolutional networks (CNNs) [11] and deep recurrent networks (RNNs) [5], have made performance breakthroughs on many real-world recognition problems. Deep architectures are very powerful for maximizing the cross-modal correlation since they can extract multilayer feature representations with different abstraction levels [1, 25, 6]. We leverage the power of deep networks for cross-modal correlation maximization by proposing a hybrid deep architecture (CHN) as illustrated in Figure 1, which constitutes an image network and a text network.




Figure 1. Correlation Hashing Network (CHN), a hybrid deep architecture for cross-modal retrieval. The image network consists of multiple convolution-pooling layers for learning image representations, and the text network consists of multiple fully-connected layers for learning text representations. Two networks are coupled by (1) a fully-connected hashing layer fch for learning hash codes, (2) a squared cosine loss for capturing both cross-modal and within-modal correlations, and (3) a cosine quantization loss for controlling quality of hash codes.

For the image network, we start with the AlexNet [11], a deep convolutional network (CNN) architecture comprised of five convolutional layers ($conv1$ – $conv5$) and three fully connected layers ($fc6$ – $fc8$). Each fc layer ℓ learns a non-linear mapping $\mathbf{u}_i^\ell = a_x^\ell(\mathbf{W}_x^\ell \mathbf{u}_i^{\ell-1} + \mathbf{b}_x^\ell)$, where \mathbf{u}_i^ℓ is the ℓ -th layer hidden representation of image \mathbf{x}_i , \mathbf{W}_x^ℓ and \mathbf{b}_x^ℓ are the weight and bias parameters of the ℓ -th layer, and a_x^ℓ is the activation function, taken as rectifier units (ReLU) $a_x^\ell(\mathbf{x}) = \max(0, \mathbf{x})$ for all hidden layers $conv1$ – $fc7$. For hash learning, we replace the $fc8$ layer (the softmax classifier layer in the original AlexNet) with a new fch layer of b hidden units, which transforms the $fc7$ representation to b -dimensional hash coding by $\mathbf{h}_i^x = \text{sgn}(\mathbf{u}_i^l)$, where $l = 8$ is the number of layers and \mathbf{u}_i^l is the hidden representation of the fch layer. To reduce the gap between the fch layer representation \mathbf{u}_i^l and the binary hash codes \mathbf{h}_i^x , we first squash the output to be within $[-1, 1]$ by utilizing the hyperbolic tangent (\tanh) activation function $a_x^l(\mathbf{x}) = \tanh(\mathbf{x})$.

For the text network, we adopt Multilayer Perceptions (MLP) [22], a widely-used architecture by the natural language processing (NLP) community. The MLP consists of three fully-connected layers, which are respectively named as $fc6$, $fc7$ and $fc8$ to keep consistent with the image network. Each fc layer ℓ learns a nonlinear mapping $\mathbf{v}_i^\ell = a_y^\ell(\mathbf{W}_y^\ell \mathbf{v}_i^{\ell-1} + \mathbf{b}_y^\ell)$, where \mathbf{v}_i^ℓ is the ℓ -th layer hidden representation of text \mathbf{y}_i , \mathbf{W}_y^ℓ and \mathbf{b}_y^ℓ are the weight and bias parameters of the ℓ -th layer, and a_y^ℓ is the activation function, taken as rectifier units (ReLU) for the hidden layer $fc6$ and $fc7$. For hash learning, we replace $fc8$ with a fch layer of b hidden units, which transforms the $fc7$ representations to b -dimensional hash coding by $\mathbf{h}_i^y = \text{sgn}(\mathbf{v}_i^l)$, where \mathbf{v}_i^l is the hidden representation of the fch layer. To reduce the gap between the fch layer representation \mathbf{v}_i^l and the binary hash codes \mathbf{h}_i^y , we also use \tanh as the activation function

of the fch layer for generating more compact hash codes. It is worth noting that, besides standard text representation methods such as tag occurrence and term-frequency inverse document frequency (TF-IDF), more advanced representation learning methods have been developed for natural language understanding, such as Word2Vec [20] and Paragraph Vector [14]. These advanced methods can be plugged into the MLP for learning more powerful text representations. In this paper, we keep our method simple and extendible by using tag occurrence as text features and MLP as the text network, while leaving advanced methods for future study.

3.2. Squared Cosine Loss

For a pair of binary codes \mathbf{h}_i^x and \mathbf{h}_j^y , there is a nice relationship between their Hamming distance $\text{dist}_H(\cdot, \cdot)$ and inner product $\langle \cdot, \cdot \rangle$: $\text{dist}_H(\mathbf{h}_i^x, \mathbf{h}_j^y) = \frac{1}{2}(b - \langle \mathbf{h}_i^x, \mathbf{h}_j^y \rangle)$. Therefore, we can employ the inner product as a reasonable surrogate of the Hamming distance to quantify the pairwise similarity. However, the approximation accuracy of such a surrogate for continuous representations \mathbf{u}_i and \mathbf{v}_j (note that $\mathbf{h}_i^x = \text{sgn}(\mathbf{u}_i)$ and $\mathbf{h}_j^y = \text{sgn}(\mathbf{v}_j)$) will be unreliable if their vector lengths are very different, i.e. $\frac{1}{2}(b - \langle \mathbf{u}_i, \mathbf{v}_j \rangle)$ will no longer be an accurate surrogate of $\text{dist}_H(\mathbf{h}_i^x, \mathbf{h}_j^y)$. Figure 2 shows such an undesirable case, where points 1 and 2 (in red) have very different vector lengths and hence very large Euclidean distance, but their Hamming distance is 0 because they are assigned with the same binary code (1, -1, 1). Such a gap between Hamming distance and inner product has raised a misspecification issue of existing inner product based deep hashing methods [16, 31, 13].

To close the gap between Hamming distance and inner product for continuous embeddings, we note that for a pair of binary codes \mathbf{h}_i^x and \mathbf{h}_j^y , there is another nice relationship between their Hamming distance $\text{dist}_H(\cdot, \cdot)$ and the cosine




Figure 2. Justifications of the proposed squared cosine loss and cosine quantization loss. (1) Similar points 1 and 2 (in red): large Euclidean distance (bad Hamming surrogate) but small cosine distance (better Hamming surrogate). (2) Similar points 3 and 4 (in purple): small cosine distance but large Hamming distance (the gap between cosine and Hamming). (3) Similar points 5 and 6 (in yellow): small cosine distance and small Hamming distance (the gap is closed by cosine quantization loss). *Best viewed in color.*

distance $\cos(\cdot, \cdot)$: $\text{dist}_H(\mathbf{h}_i^x, \mathbf{h}_j^y) = \frac{b}{2} (1 - \cos(\mathbf{h}_i^x, \mathbf{h}_j^y))$, where $\cos(\mathbf{u}_i, \mathbf{v}_j) = \frac{\langle \mathbf{u}_i, \mathbf{v}_j \rangle}{\|\mathbf{u}_i\| \|\mathbf{v}_j\|}$, and $\|\cdot\|$ is the vector length. Since cosine distance can mitigate the diversity of vector lengths and make the continuous representations \mathbf{u}_i and \mathbf{v}_j lie on the unit sphere (only the angle between them is measured), it makes $\frac{b}{2} (1 - \cos(\mathbf{u}_i, \mathbf{v}_j))$ a much more accurate surrogate of $\text{dist}_H(\mathbf{h}_i^x, \mathbf{h}_j^y)$. As can be seen in Figure 2, the cosine distance between points 1 and 2 (in red) is closer to zero and thus better approximates their Hamming distance. Hence in this paper, we opt to employ the cosine distance as a good surrogate of the Hamming distance, which leads to a novel squared cosine loss for correlation maximization.

3.2.1 Cross-Modal Correlation

To maximize the cross-modal correlation, we propose the following criterion: for each pair of objects $(\mathbf{o}_i, \mathbf{o}_j, s_{ij})$, if $s_{ij} = 1$, indicating that \mathbf{o}_i and \mathbf{o}_j are similar, then their binary hash codes must be similar across different modalities (image and text), i.e. the Hamming distance should satisfy $d_H(\mathbf{h}_i^x, \mathbf{h}_j^y) \rightarrow 0$ and $d_H(\mathbf{h}_i^y, \mathbf{h}_j^x) \rightarrow 0$, which implies the cosine distance should satisfy $\cos(\mathbf{u}_i, \mathbf{v}_j) \rightarrow 1$ and $\cos(\mathbf{v}_i, \mathbf{u}_j) \rightarrow 1$. Correspondingly, if $s_{ij} = -1$, indicating that \mathbf{o}_i and \mathbf{o}_j are dissimilar, then by derivation, the cosine distance should satisfy $\cos(\mathbf{u}_i, \mathbf{v}_j) \rightarrow -1$ and $\cos(\mathbf{v}_i, \mathbf{u}_j) \rightarrow -1$. It is very important to note that, for other distance metrics (e.g. inner product, Euclidean distance, etc), it is very difficult to devise such a well-specified learning criterion because these distances are not good surrogates of the Hamming distance. The above learning criterion leads to a novel squared cosine

loss for maximizing the cross-modal correlation as follows,

$$C_{xy} = \sum_{s_{ij} \in \mathcal{S}} (s_{ij} - \cos(\mathbf{u}_i, \mathbf{v}_j))^2 + (s_{ij} - \cos(\mathbf{v}_i, \mathbf{u}_j))^2, \quad (1)$$

The range of cosine distance $\cos(\mathbf{u}_i, \mathbf{v}_j) \in [-1, 1]$ is consistent with binary similarity labels $s_{ij} \in \{-1, 1\}$, making the squared cosine loss in Equation (1) a well-specified loss for preserving the pairwise similarity information conveyed in \mathcal{S} . The squared cosine loss is particularly powerful for cross-modal correlation analysis, since the vector lengths are very diverse in different modalities and may make other distance metrics (e.g. Euclidean distance) misspecified. In real retrieval systems, cosine distance is widely adopted to mitigate the diversity of vector lengths and significantly improve the retrieval quality, but it has not been used in deep learning to hash methods for the cross-modal retrieval [27].

3.2.2 Within-Modal Correlation

Only maximizing the cross-modal correlation is not sufficient for learning high-quality hash codes for cross-modal retrieval. The justifications are two folds: (1) Theoretically, each modality may have its modality-specific structure (image or text), which cannot be well captured by exploring its correlation with another modality; (2) Technically, deep neural networks cannot receive sufficient training if within-modal knowledge structure is not explored (due to gradient vanishing issue). Hence, we further maximize the within-modal correlation using the squared cosine loss as follows,

$$\begin{aligned} C_{xx} &= \sum_{s_{ij} \in \mathcal{S}} (s_{ij} - \cos(\mathbf{u}_i, \mathbf{u}_j))^2, \\ C_{yy} &= \sum_{s_{ij} \in \mathcal{S}} (s_{ij} - \cos(\mathbf{v}_i, \mathbf{v}_j))^2. \end{aligned} \quad (2)$$

By maximizing the within-modal correlation, we can learn high-quality deep representations for both image and text by using the image network and text network, respectively. The high-quality modality-specific deep representations guarantee that the cross-modal correlation can be captured and enhanced by the hybrid deep architecture while the different convergence rates between the image network (CNNs) and the text network (MLPs) will not cause a training difficulty.

3.3. Cosine Quantization Loss

Although we justify that the cosine distance is a good surrogate of the Hamming distance by $\text{dist}_H(\mathbf{h}_i^x, \mathbf{h}_j^y) \approx \frac{b}{2} (1 - \cos(\mathbf{u}_i, \mathbf{v}_j))$, such an approximation may fail when two similar points \mathbf{u}_i and \mathbf{v}_j with $s_{ij} = 1$ (i.e. their cosine distance is small due to minimizing the squared cosine loss) lie on different sides of the hyperplane (i.e. their Hamming distance is large due to different signs of hash codes across the hyperplane). Figure 2 shows such an undesirable case,

where points 3 and 4 (in purple) have small cosine distance but large Hamming distance because they are assigned with different binary codes $(1, -1, 1)$ and $(1, 1, 1)$, respectively. Such a gap between Hamming distance and cosine distance may result in an inaccurate surrogate approximation.

To close the gap between Hamming distance and cosine distance for continuous embeddings, we note that for a pair of continuous representations \mathbf{u}_i and \mathbf{v}_j , if they are close (in cosine distance) to their signed codes $\mathbf{h}_i^x = \text{sgn}(\mathbf{u}_i)$ and $\mathbf{h}_i^y = \text{sgn}(\mathbf{v}_i)$ (i.e. far from the hyperplane), then they will lie in the same hypercube with high probability (i.e. have the same binary code and hence their Hamming distance is zero). Figure 2 shows such a desirable case by points 5 and 6 (in yellow). Hence we need to maximize $\cos(\mathbf{u}_i, \text{sgn}(\mathbf{u}_i))$ and $\cos(\mathbf{v}_j, \text{sgn}(\mathbf{v}_j))$, which is equivalent to maximizing $\cos(|\mathbf{u}_i|, \mathbf{1})$ and $\cos(|\mathbf{v}_j|, \mathbf{1})$, respectively. The observation leads to a novel cosine quantization loss defined as follows,

$$\begin{aligned} Q_x &= - \sum_{s_{ij} \in \mathcal{S}} (\cos(|\mathbf{u}_i|, \mathbf{1}) + \cos(|\mathbf{u}_j|, \mathbf{1})), \\ Q_y &= - \sum_{s_{ij} \in \mathcal{S}} (\cos(|\mathbf{v}_i|, \mathbf{1}) + \cos(|\mathbf{v}_j|, \mathbf{1})), \end{aligned} \quad (3)$$

where $\mathbf{1} \in \mathbb{R}^b$ is the vector of ones. Note that, minimizing the cosine quantization loss will not only close the gap between the Hamming distance and cosine distance, but also lead to lower quantization error when binarizing the continuous representations $\mathbf{u}_i \in \mathbb{R}^b$ and $\mathbf{v}_j \in \mathbb{R}^b$ to hash codes $\mathbf{h}_i^x = \text{sgn}(\mathbf{u}_i) \in \{-1, 1\}^b$ and $\mathbf{h}_j^y = \text{sgn}(\mathbf{v}_j) \in \{1, -1\}^b$. This advantage enables us to learn high-quality hash codes.

3.4. Joint Optimization Problem

We perform simultaneous representation learning and hash coding by jointly maximizing the cross-modal/within-modal correlations and controlling the quantization error. Integrating (1)–(3) into a joint optimization problem yields

$$\min_{\Theta} L = C_{xy} + \lambda(C_{xx} + C_{yy}) + \gamma(Q_x + Q_y), \quad (4)$$

where $\Theta \triangleq \{\mathbf{W}^\ell, \mathbf{b}^\ell\}$ is the set of network parameters, λ is the penalty parameter to control the relative importance of within-modal squared cosine loss, and γ is the penalty parameter for trading off the relative importance of the cosine quantization loss. Through joint optimization problem (4), we can achieve optimal learning of compact hash codes for cross-modal retrieval, e.g. text query on image database. Finally, we can obtain b -bit binary codes by simple quantization $\mathbf{h}_x \leftarrow \text{sgn}(\mathbf{u})$ and $\mathbf{h}_y \leftarrow \text{sgn}(\mathbf{v})$, where $\text{sgn}(\mathbf{z})$ is the sign function on vectors that for $i = 1, \dots, b$, $\text{sgn}(z_i) = 1$ if $z_i > 0$, otherwise $\text{sgn}(z_i) = -1$. It is worth noting that, since we have minimized the cosine quantization error in (4) during training, this final binarization step will incur very small loss of retrieval quality as validated empirically.

4. Algorithm and Analysis

4.1. Learning Algorithm

We derive the learning algorithms for the CHN model in Equation (4), and show rigorously that both squared cosine loss and cosine quantization loss can be optimized efficiently through the standard back-propagation (BP) procedure. For notation brevity, we define the pointwise cost as

$$\begin{aligned} L_i &\triangleq \sum_{j:s_{ij} \in \mathcal{S}} C_{ij}^{xy} + \lambda(C_{ij}^{xx} + C_{ij}^{yy}) + \gamma(Q_{ij}^x + Q_{ij}^y) \\ &= \sum_{j:s_{ij} \in \mathcal{S}} \left(s_{ij} - \cos(\mathbf{u}_i^l, \mathbf{v}_j^l) \right)^2 + \left(s_{ij} - \cos(\mathbf{v}_i^l, \mathbf{u}_j^l) \right)^2 \\ &\quad + \lambda \sum_{j:s_{ij} \in \mathcal{S}} \left(s_{ij} - \cos(\mathbf{u}_i^l, \mathbf{u}_j^l) \right)^2 + \left(s_{ij} - \cos(\mathbf{v}_i^l, \mathbf{v}_j^l) \right)^2 \\ &\quad - \gamma \sum_{j:s_{ij} \in \mathcal{S}} \left(\left(\cos(|\mathbf{u}_i^l|, \mathbf{1}) + \cos(|\mathbf{u}_j^l|, \mathbf{1}) \right) \right) \\ &\quad - \gamma \sum_{j:s_{ij} \in \mathcal{S}} \left(\left(\cos(|\mathbf{v}_i^l|, \mathbf{1}) + \cos(|\mathbf{v}_j^l|, \mathbf{1}) \right) \right), \end{aligned} \quad (5)$$

where $\mathbf{u}_i = \mathbf{u}_i^l$ and $\mathbf{v}_j = \mathbf{v}_j^l$ denote the continuous representations of the hybrid network. We derive the gradient of point-wise cost L_i w.r.t. $\mathbf{W}_{x,k}^\ell$, the network parameter of the k -th unit in the ℓ -th layer for the image network (the derivation of the text network is the same and is omitted) as

$$\begin{aligned} \frac{\partial L_i}{\partial \mathbf{W}_{x,k}^\ell} &= \sum_{j:s_{ij} \in \mathcal{S}} \left(\frac{\partial C_{ij}^{xy}}{\partial \mathbf{W}_{x,k}^\ell} + \lambda \frac{\partial C_{ij}^{xx} + C_{ij}^{yy}}{\partial \mathbf{W}_{x,k}^\ell} \right) \\ &\quad + \sum_{j:s_{ij} \in \mathcal{S}} \left(\gamma \frac{\partial Q_{ij}^x + Q_{ij}^y}{\partial \mathbf{W}_{x,k}^\ell} \right) \\ &= \sum_{j:s_{ij} \in \mathcal{S}} \left(\frac{\partial C_{ij}^{xy}}{\partial \mathbf{W}_{x,k}^\ell} + \lambda \frac{\partial C_{ij}^{xx}}{\partial \mathbf{W}_{x,k}^\ell} + \gamma \frac{\partial Q_{ij}^x}{\partial \mathbf{W}_{x,k}^\ell} \right) \\ &= \sum_{j:s_{ij} \in \mathcal{S}} \left(\frac{\partial C_{ij}^{xy}}{\partial \hat{\mathbf{u}}_{ik}^\ell} + \lambda \frac{\partial C_{ij}^{xx}}{\partial \hat{\mathbf{u}}_{ik}^\ell} + \gamma \frac{\partial Q_{ij}^x}{\partial \hat{\mathbf{u}}_{ik}^\ell} \right) \frac{\partial \hat{\mathbf{u}}_{ik}^\ell}{\partial \mathbf{W}_{x,k}^\ell} \\ &= \delta_{x,ik}^\ell \mathbf{u}_i^{\ell-1}, \end{aligned} \quad (6)$$

where $\hat{\mathbf{u}}_i^\ell = \mathbf{W}_x^\ell \mathbf{u}_i^{\ell-1} + \mathbf{b}_x^\ell$ is ℓ -th layer output before activation $a_x^\ell(\cdot)$, $\delta_{x,ik}^\ell \triangleq \sum_{j:s_{ij} \in \mathcal{S}} \left(\frac{\partial C_{ij}^{xy}}{\partial \hat{\mathbf{u}}_{ik}^\ell} + \lambda \frac{\partial C_{ij}^{xx}}{\partial \hat{\mathbf{u}}_{ik}^\ell} + \gamma \frac{\partial Q_{ij}^x}{\partial \hat{\mathbf{u}}_{ik}^\ell} \right)$ is

the point-wise *residual* term that measures how much the k -th unit in the ℓ -th layer is responsible for the error of point \mathbf{x}_i in the network output. For an output unit k , we can measure the difference between the network's activation and the true target value, and use that to define the residual $\delta_{x,ik}^\ell$ as

$$\begin{aligned} \delta_{x,ik}^\ell &= \left(\sum_{j:s_{ij} \in \mathcal{S}} \left[\left(\frac{\langle \mathbf{u}_i^l, \mathbf{v}_j^l \rangle}{\|\mathbf{u}_i^l\| \|\mathbf{v}_j^l\|} - s_{ij} \right) \left(\frac{\mathbf{v}_{jk}^l}{\|\mathbf{u}_i^l\| \|\mathbf{v}_j^l\|} - \frac{\mathbf{u}_{ik}^l \langle \mathbf{u}_i^l, \mathbf{v}_j^l \rangle}{\|\mathbf{u}_i^l\|^3 \|\mathbf{v}_j^l\|} \right) \right] \right. \\ &\quad + \lambda \sum_{j \neq i:s_{ij} \in \mathcal{S}} \left[\left(\frac{\langle \mathbf{u}_i^l, \mathbf{u}_j^l \rangle}{\|\mathbf{u}_i^l\| \|\mathbf{u}_j^l\|} - s_{ij} \right) \left(\frac{\mathbf{u}_{jk}^l}{\|\mathbf{u}_i^l\| \|\mathbf{u}_j^l\|} - \frac{\mathbf{u}_{ik}^l \langle \mathbf{u}_i^l, \mathbf{u}_j^l \rangle}{\|\mathbf{u}_i^l\|^3 \|\mathbf{u}_j^l\|} \right) \right] \\ &\quad \left. + \gamma \sum_{j:s_{ij} \in \mathcal{S}} \left[\frac{\text{sgn}(u_{ik})}{\sqrt{b} \|\mathbf{u}_i\|} - \frac{u_{ik} \langle \mathbf{u}_i, \mathbf{1} \rangle}{\sqrt{b} \|\mathbf{u}_i\|^3} \right] \right) \cdot 2a^l(\hat{\mathbf{u}}_{ik}^\ell) \end{aligned} \quad (7)$$

where $l = 8$ denotes the index of the output layer, and $\dot{a}_x(\cdot)$ is the derivative of the l -th layer activation function. For a hidden unit k in the $(\ell - 1)$ -th layer, we compute the residual $\delta_{x,ik}^{\ell-1}$ based on a weighted average of the errors of all the units $k' = 1, \dots, n_{\ell-1}$ in the $(\ell - 1)$ -th layer that use $\mathbf{u}_i^{\ell-1}$ as an input, which is just consistent with the BP procedure,

$$\delta_{x,ik}^{\ell-1} = \left(\sum_{k'=1}^{n_{\ell-1}} \delta_{x,ik'}^{\ell} W_{x,kk'}^{\ell-1} \right) \dot{a}_x^{\ell-1}(\hat{u}_{ik}^{\ell-1}), \quad (8)$$

where $n_{\ell-1}$ is number of units in $(\ell - 1)$ -th layer. The residuals in all layers can be computed by back-propagation.

An important property of the proposed algorithm is that, only computing the residual of the output layer involves the pairwise summation as in Equation (7). For all hidden layers, all the residuals can be simply computed recursively by Equation (8), which does not involve pairwise summation. Hence we do not need to modify the implementation of BP in all hidden layers $1 \leq \ell \leq l - 1$. We only modify standard BP by replacing the output residual with Equation (7).

Since the only difference between standard BP and our algorithm is Equation (7), we analyze the computational complexity based on Equation (7). Denote the number of similarity pairs \mathcal{S} available for training as $|\mathcal{S}|$, then it is easy to verify that the computational complexity is linear $O(|\mathcal{S}|)$.

4.2. Theoretical Analysis

We elaborate the connection between our work and Iterative Quantization (ITQ) [7], the seminal work that considers the quantization error in an iterative Procrustean procedure. In ITQ, point-wise quantization error is defined as follows,

$$Q_{\text{ITQ}} = \|\mathbf{u}_i - \text{sgn}(\mathbf{u}_i)\|^2. \quad (9)$$

We only consider point-wise quantization error for brevity, while taking the summation for all training points (pairs) yields the overall quantization error similar to Equation (3).

Theorem 1 (Upper Bound). *The proposed cosine quantization loss in (3) is proportional to an upper bound of the ITQ quantization error in (9) (up to a constant $2b$),*

$$\|\mathbf{u}_i - \text{sgn}(\mathbf{u}_i)\|^2 \leq 2b - 2b \cos(|\mathbf{u}_i|, \mathbf{1}). \quad (10)$$

Proof. It follows from triangle inequality $\||\mathbf{u}_i| - \mathbf{1}\|^2 = \||\mathbf{u}_i| - |\text{sgn}(\mathbf{u}_i)|\|^2 \leq \|\mathbf{u}_i - \text{sgn}(\mathbf{u}_i)\|^2$. Since \mathbf{u}_i and $\text{sgn}(\mathbf{u}_i)$ have the same sign, the last equality always holds, hence $\|\mathbf{u}_i - \text{sgn}(\mathbf{u}_i)\|^2 = \||\mathbf{u}_i| - \mathbf{1}\|^2$. Because we use hyperbolic tangent (\tanh) as the activation function for the

fch layer, $|\mathbf{u}_i| \leq 1$, $\|\mathbf{u}_i\|^2 \leq b$, and $\langle |\mathbf{u}_i|, \mathbf{1} \rangle \geq 0$. Hence,

$$\begin{aligned} \||\mathbf{u}_i| - \mathbf{1}\|^2 &= \|\mathbf{u}_i\|^2 \left\| \frac{|\mathbf{u}_i|}{\|\mathbf{u}_i\|} - \frac{\mathbf{1}}{\|\mathbf{u}_i\|} \right\|^2 \\ &\leq b \left\| \frac{|\mathbf{u}_i|}{\|\mathbf{u}_i\|} - \frac{\mathbf{1}}{\|\mathbf{u}_i\|} \right\|^2 = 2b \left(1 - \frac{\langle |\mathbf{u}_i|, \mathbf{1} \rangle}{\|\mathbf{u}_i\|^2} \right) \\ &\leq 2b \left(1 - \frac{\langle |\mathbf{u}_i|, \mathbf{1} \rangle}{\|\mathbf{u}_i\| \sqrt{b}} \right) = 2b - 2b \cos(|\mathbf{u}_i|, \mathbf{1}). \end{aligned} \quad (11)$$

□

Theorem 1 reveals that the cosine quantization loss (3) is a reasonable criterion for hash learning. Different from all existing methods, the proposed cosine quantization loss is jointly optimized with the squared cosine loss in a hybrid deep network, which can yield better hashing schemes by jointly (1) reducing the gap between Hamming distance and cosine distance and (2) controlling the binary coding error.

5. Experiments

We conduct extensive experiments to evaluate the efficacy of the proposed CHN model with several state of the art hashing methods on widely-used benchmark datasets. The codes and configurations will be available online.

5.1. Evaluation Setup

We conduct our extensive evaluation on two benchmark cross-modal datasets: **NUS-WIDE** [4] and **MIR-Flickr** [9].

NUS-WIDE¹ is a public web image dataset containing 269,648 images downloaded from Flickr, together with the associated raw tags of these images. There are 81 ground truth concepts manually annotated for search evaluation. Following prior works [36, 28], we prune the original NUS-WIDE dataset to form a new dataset consisting of 195,834 image-text pairs that belong to one of the 21 most frequent concepts. All images are resized into 256×256 .

MIR-Flickr includes 1 million images associated with tags from Flickr, where 25,000 are labeled with 38 concepts. We resize images of this labeled subset into 256×256 .

For our deep learning based approach CHN, we directly use the raw image pixels as the input. For the shallow learning based methods, we follow [16, 13, 25] to represent each image in NUS-WIDE by a 500-dimensional bag-of-visual-words vector, and to represent each image in Flickr by a 3,857-dimensional vector concatenated by local SIFT feature, global GIST feature, and etc. For text modality, all the methods use tag occurrence vectors as the input. All image and text features are available at the datasets' website. In NUS-WIDE, we randomly select 100 pairs per class as the test query set, 500 pairs per class as the training set and

¹<http://lms.comp.nus.edu.sg/research/NUS-WIDE.htm>

Table 1. Mean Average Precision (MAP) Comparison of Two Cross-modal Retrieval Tasks on Two Datasets

Task	Method	NUS-WIDE				MIR-Flickr			
		8 bits	16 bits	32 bits	64 bits	8 bits	16 bits	32 bits	64 bits
$I \rightarrow T$	CMSSH	0.3950	0.4052	0.4076	0.3516	0.5076	0.5272	0.5357	0.5219
	CVH	0.4588	0.4713	0.4743	0.4740	0.6091	0.6225	0.6364	0.6199
	IMH	0.4345	0.4399	0.4203	0.4115	0.5449	0.5646	0.5936	0.5539
	SCM	0.4693	0.4648	0.4619	0.4851	0.6361	0.6493	0.6495	0.6440
	QCH	0.4765	0.4895	0.5050	0.5125	0.6452	0.6523	0.6685	0.6721
	CorrAE	0.4398	0.4522	0.4699	0.4944	0.6301	0.6329	0.6357	0.6401
	CM-NN	<u>0.6035</u>	<u>0.6255</u>	<u>0.6424</u>	<u>0.6514</u>	<u>0.7154</u>	<u>0.7288</u>	<u>0.7415</u>	<u>0.7541</u>
	CHN	0.7126	0.7692	0.8220	0.8277	0.7890	0.8430	0.8571	0.8572
$T \rightarrow I$	CMSSH	0.3783	0.3499	0.3944	0.4015	0.5868	0.5732	0.6176	0.6323
	CVH	0.5598	0.5217	0.5129	0.4875	0.5972	0.6032	0.5738	0.5794
	IMH	0.4380	0.4582	0.4186	0.4051	0.5374	0.5536	0.5513	0.5583
	SCM	0.4449	0.4859	0.5105	0.5259	0.6037	0.5998	0.5805	0.6078
	QCH	0.5020	0.5195	0.5489	0.5622	0.6258	0.6425	0.6485	0.6528
	CorrAE	0.4303	0.4501	0.4634	0.4880	0.6142	0.6198	0.6247	0.6431
	CM-NN	<u>0.6012</u>	<u>0.6083</u>	<u>0.6226</u>	<u>0.6435</u>	<u>0.6969</u>	<u>0.7142</u>	<u>0.7285</u>	<u>0.7352</u>
	CHN	0.7170	0.7617	0.7704	0.7649	0.7595	0.7631	0.7814	0.7886




Figure 3. Precision-recall curves of the two cross-modal retrieval tasks on the NUS-WIDE and Flickr datasets with 32 bits hash codes.

50 pairs per class as the validation set. In Flickr, we randomly select 1000 pairs as the test query set, 4000 pairs as the training set and 1000 pairs as the validation set. The similarity pairs for training are randomly constructed using semantic labels: each pair is considered similar (dissimilar) if they share at least one (none) semantic label.

We evaluate and compare the retrieval accuracies of the proposed CHN with 7 state-of-the-art cross-modal hashing methods, including three unsupervised methods **IMH**² [24], **CVH**³ [12] and **CorrAE**⁴ [6], and four supervised methods **CMSSH**³ [3], **SCM**⁴ [33], **QCH**⁴ [30] and **CM-NN**⁴ [19].

We follow [19, 6, 33, 30] to evaluate the retrieval quality based on three widely used metrics: Mean Average Precision (MAP), *precision-recall* curves [36, 28] and *precision@top-R* curves. We adopt MAP@R = 50 following previous work [18, 32, 28].

We implement the CHN model based on the open-source **Caffe** framework [10]. For image network, we employ the AlexNet architecture [11], fine-tune convolutional layers *conv1-conv5* and fully-connected layers *fc6-fc7* that were copied from the pre-trained model, and train hash-

ing layer *fch*, all via back-propagation. As the *fch* layer is trained from scratch, we set its learning rate to be 10 times that of the lower layers. For text network, we employ a three layer multi-layer perceptrons (MLP), in which the *fc6-fc7* layers have 256 ReLU units with dropout rate 0.5, and the *fch* layer have *b* tanh units. We use the mini-batch stochastic gradient descent (SGD) with 0.9 momentum and the learning rate annealing strategy implemented in Caffe, cross-validate learning rate from 10^{-5} to 1 with a multiplicative step-size 10, and fix mini-batch size as 64.

The CHN approach involves two model parameters, coefficient λ and γ , the relative weights of the within-modal cosine loss and the cosine quantization loss. We automatically select λ and γ using cross-validation. For the comparison methods, we use cross-validation to carefully tune their parameters, respectively. Each experiment repeats ten runs and the average results are reported.

5.2. Results and Discussions

We compare CHN with seven state-of-the-art methods in terms of MAP, precision-recall and precision@top-R curves of two cross-modal retrieval tasks: image query on text database ($I \rightarrow T$), text query on image database ($T \rightarrow I$).

²http://staff.itee.uq.edu.au/shenht/UQ_IMH

³<http://www.cse.ust.hk/~dyyeung/code/mlbe.zip>

⁴Since the original code is not publicly available, we implement it by ourselves.

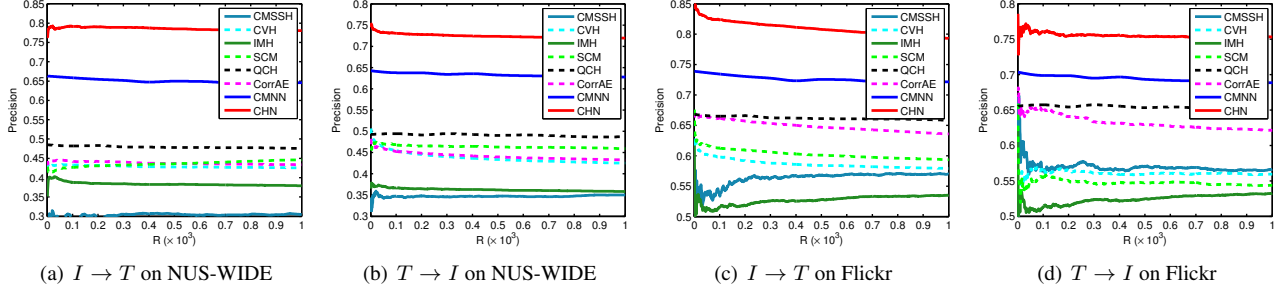


Figure 4. Precision@top R curves of the two cross-modal retrieval tasks on the NUS-WIDE and Flickr datasets with 32 bits hash codes.

Table 2. Mean Average Precision (MAP) Comparison of Two Cross-modal Retrieval Tasks of CHN variants on Two Datasets.

Task	Method	NUS-WIDE				MIR-Flickr			
		8 bits	16 bits	32 bits	64 bits	8 bits	16 bits	32 bits	64 bits
$I \rightarrow T$	CHN-W	0.3257	0.3718	0.3696	0.3580	0.5593	0.5676	0.5810	0.6069
	CHN-C	0.6214	0.6748	0.7235	0.7582	0.7584	0.7878	0.7942	0.7992
	CHN-Q	<u>0.6852</u>	<u>0.7445</u>	<u>0.7890</u>	<u>0.8035</u>	<u>0.7671</u>	<u>0.8056</u>	<u>0.8173</u>	<u>0.8229</u>
	CHN	0.7126	0.7692	0.8220	0.8277	0.7890	0.8430	0.8571	0.8572
$T \rightarrow I$	CHN-W	0.3368	0.3500	0.3262	0.3868	0.4346	0.4649	0.4906	0.5203
	CHN-C	0.6045	0.6327	0.6514	0.6790	0.7215	0.7349	0.7396	0.7581
	CHN-Q	<u>0.6922</u>	<u>0.7418</u>	<u>0.7562</u>	<u>0.7581</u>	<u>0.7384</u>	<u>0.7519</u>	<u>0.7667</u>	<u>0.7784</u>
	CHN	0.7170	0.7617	0.7704	0.7649	0.7595	0.7631	0.7814	0.7886

We evaluate all methods with different lengths of hash codes, i.e. 8, 16, 32, and 64 bits, and report their MAP results in Table 1. From the experimental results, we can observe that CHN substantially outperforms all state-of-the-art methods for all cross-modal tasks on the two benchmark datasets. Specifically, compared to the best shallow baseline QCH, we achieve absolute increases of 28.70% / 22.04% and 17.71% / 13.08% in average MAP for two cross-modal tasks $I \rightarrow T$ / $T \rightarrow I$ on NUS-WIDE and Flickr, respectively. CHN outperforms the state-of-the-art deep hashing method CM-NN by 15.22% / 13.46% and 10.16% / 5.45% in average MAP for two cross-modal tasks on the datasets.

The precision-recall curves with 32 bits for two cross-modal tasks $I \rightarrow T$ and $T \rightarrow I$ on two benchmark datasets NUS-WIDE and Flickr are illustrated in Figure 3, respectively. CHN shows the best cross-modal retrieval performance at all recall levels. This validates that CHN is robust to diverse retrieval scenarios, such as higher recall tolerating lower precision. Figure 4 shows the precision@top-R curves of all state-of-the-art methods, which further represent the precision changes along with the number of top retrieved results R with 32 bits on NUS-WIDE and Flickr datasets, respectively. We can easily observe that CHN significantly outperforms all of the state-of-the-art methods, which shows that CHN is suitable for the applications that prefer higher precision with fewer top-R retrieved results.

5.3. Empirical Analysis

In order to study the effectiveness of our approach, we investigate several variants of CHN: CHN-W is the CHN

variant with only within-modal squared cosine loss C_{xx} and C_{yy} (2); CHN-C is the CHN variant with only the cross-modal squared cosine loss C_{xy} (1); CHN-Q is a stronger baseline with both within-modal and cross-modal squared cosine loss, but without the cosine quantization loss (3). We report the MAP results of all CHN variants in Table 2.

From Table 2, we may have the following observations: (1) CHN-Q outperforms CHN-W for very large margins of 39.93% / 38.71% and 22.45% / 28.13% in average MAP on two cross-modal tasks on two benchmark datasets, confirming the vital importance of maximizing the cross-modal correlation. (2) CHN-Q also significantly outperforms CHN-C by 6.11% / 9.52% and 1.83% / 2.03% in average MAP, confirming that by preserving the within-modal correlation, the network can better exploit modality-specific knowledge structure for learning more powerful deep representations. (3) Compared to CHN-Q, the MAP results of CHN demonstrate that the cosine quantization loss can evidently reduce the loss when binarizing the continuous representations to hash codes and achieve better cross-modal retrieval results.

6. Conclusion

In this paper, we proposed a Correlation Hashing Network (CHN) for cross-modal retrieval. CHN is a hybrid architecture that jointly optimizes the squared cosine loss on semantic similarity pairs and the cosine quantization loss on compact hash codes. Extensive experiments on standard cross-media retrieval datasets show the CHN model yields substantial boosts over the state-of-the-art hashing methods.

References

- [1] G. Andrew, R. Arora, J. Bilmes, and K. Livescu. Deep canonical correlation analysis. In *Proc. 30th ICML*, 2013. 2
- [2] Y. Bengio, A. Courville, and P. Vincent. Representation learning: A review and new perspectives. *TPAMI*, 35, 2013. 2
- [3] M. Bronstein, A. Bronstein, F. Michel, and N. Paragios. Data fusion through cross-modality metric learning using similarity-sensitive hashing. In *CVPR*. IEEE, 2010. 1, 2, 7
- [4] T.-S. Chua, J. Tang, R. Hong, H. Li, Z. Luo, and Y.-T. Zheng. Nus-wide: A real-world web image database from national university of singapore. In *CIVR*. ACM, 2009. 6
- [5] J. Donahue, L. A. Hendricks, S. Guadarrama, M. Rohrbach, S. Venugopalan, T. Darrell, and K. Saenko. Long-term recurrent convolutional networks for visual recognition and description. In *CVPR 2015*, 2015. 2
- [6] F. Feng, X. Wang, and R. Li. Cross-modal retrieval with correspondence autoencoder. In *MM*. ACM, 2014. 1, 2, 7
- [7] Y. Gong and S. Lazebnik. Iterative quantization: A procrustean approach to learning binary codes. In *CVPR 2011*. IEEE, 2011. 6
- [8] Y. Hu, Z. Jin, H. Ren, D. Cai, and X. He. Iterative multi-view hashing for cross media indexing. In *MM*. ACM, 2014. 2
- [9] M. J. Huiskes and M. S. Lew. The mir flickr retrieval evaluation. In *ICMR*. ACM, 2008. 6
- [10] Y. Jia, E. Shelhamer, J. Donahue, S. Karayev, J. Long, R. Girshick, S. Guadarrama, and T. Darrell. Caffe: Convolutional architecture for fast feature embedding. In *ACM Multimedia Conference*. ACM, 2014. 7
- [11] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in Neural Information Processing Systems (NIPS)*, 2012. 1, 2, 3, 7
- [12] S. Kumar and R. Udupa. Learning hash functions for cross-view similarity search. In *IJCAI*, 2011. 1, 2, 7
- [13] H. Lai, Y. Pan, Y. Liu, and S. Yan. Simultaneous feature learning and hash coding with deep neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, 2015. 1, 2, 3, 6
- [14] Q. V. Le and T. Mikolov. Distributed representations of sentences and documents. In *Advances in neural information processing systems*, 2014. 3
- [15] M. Lin, Q. Chen, and S. Yan. Network in network. In *International Conference on Learning Representations (ICLR)*, 2014 (arXiv:1409.1556), 2014. 1
- [16] W. Liu, J. Wang, R. Ji, Y.-G. Jiang, and S.-F. Chang. Supervised hashing with kernels. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, 2012. 3, 6
- [17] X. Liu, J. He, C. Deng, and B. Lang. Collaborative hashing. In *CVPR*. IEEE, 2014. 1, 2
- [18] X. Lu, F. Wu, S. Tang, Z. Zhang, X. He, and Y. Zhuang. A low rank structural large margin method for cross-modal ranking. In *SIGIR*. ACM, 2013. 7
- [19] J. Masci, M. M. Bronstein, A. M. Bronstein, and J. Schmidhuber. Multimodal similarity-preserving hashing. *IEEE Trans. Pattern Anal. Mach. Intell.*, 36, 2014. 1, 2, 7
- [20] T. Mikolov, I. Sutskever, K. Chen, G. S. Corrado, and J. Dean. Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*, 2013. 3
- [21] J. C. Pereira, E. Coviello, G. Doyle, N. Rasiwasia, G. R. G. Lanckriet, R. Levy, and N. Vasconcelos. On the role of correlation and abstraction in cross-modal multimedia retrieval. *TPAMI*, 36, 2014. 1
- [22] D. E. Rumelhart, G. E. Hinton, and R. J. Williams. Parallel distributed processing: Explorations in the microstructure of cognition, vol. 1. chapter Learning Internal Representations by Error Propagation. MIT Press, 1986. 3
- [23] A. W. Smeulders, M. Worring, S. Santini, A. Gupta, and R. Jain. Content-based image retrieval at the end of the early years. *TPAMI*, 22, 2000. 1, 2
- [24] J. Song, Y. Yang, Y. Yang, Z. Huang, and H. T. Shen. Inter-media hashing for large-scale retrieval from heterogeneous data sources. In *SIGMOD*. ACM, 2013. 1, 2, 7
- [25] N. Srivastava and R. Salakhutdinov. Multimodal learning with deep boltzmann machines. *JMLR*, 15, 2014. 1, 2, 6
- [26] J. Wan, D. Wang, S. C. H. Hoi, P. Wu, J. Zhu, Y. Zhang, and J. Li. Deep learning for content-based image retrieval: A comprehensive study. In *MM*. ACM, 2014. 1
- [27] J. Wang, H. T. Shen, J. Song, and J. Ji. Hashing for similarity search: A survey. Arxiv, 2014. 1, 2, 4
- [28] W. Wang, B. C. Ooi, X. Yang, D. Zhang, and Y. Zhuang. Effective multi-modal retrieval based on stacked auto-encoders. In *VLDB*. ACM, 2014. 1, 2, 6, 7
- [29] Y. Weiss, A. Torralba, and R. Fergus. Spectral hashing. In *NIPS*, 2009. 2
- [30] B. Wu, Q. Yang, W. Zheng, Y. Wang, and J. Wang. Quantized correlation hashing for fast cross-modal search. In *Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence, IJCAI 2015, Buenos Aires, Argentina, July 25-31, 2015*, 2015. 1, 2, 7
- [31] R. Xia, Y. Pan, H. Lai, C. Liu, and S. Yan. Supervised hashing for image retrieval via image representation learning. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*. AAAI, 2014. 1, 2, 3
- [32] Z. Yu, F. Wu, Y. Yang, Q. Tian, J. Luo, and Y. Zhuang. Discriminative coupled dictionary hashing for fast cross-media retrieval. In *SIGIR*. ACM, 2014. 1, 2, 7
- [33] D. Zhang and W. Li. Large-scale supervised multimodal hashing with semantic correlation maximization. In *AAAI*, 2014. 1, 2, 7
- [34] Y. Zhen and D.-Y. Yeung. Co-regularized hashing for multimodal data. In *NIPS*, 2012. 1, 2
- [35] Y. Zhen and D.-Y. Yeung. A probabilistic model for multimodal hash function learning. In *SIGKDD*. ACM, 2012. 1, 2
- [36] X. Zhu, Z. Huang, H. T. Shen, and X. Zhao. Linear cross-modal hashing for efficient multimedia search. In *MM*. ACM, 2013. 6, 7